

# “Translation can’t change a name”: Using Multilingual Data for Named Entity Recognition

Manaal Faruqui

Language Technologies Institute  
Carnegie Mellon University  
Pittsburgh, PA, 15213, USA  
mfaruqui@cs.cmu.edu

## Abstract

Named Entities (NEs) are often written with no orthographic changes across different languages that share a common alphabet. We show that this can be leveraged so as to improve named entity recognition (NER) by using unsupervised word clusters from secondary languages as features in state-of-the-art discriminative NER systems. We observe significant increases in performance, finding that person and location identification is particularly improved, and that phylogenetically close languages provide more valuable features than more distant languages.

## 1 Introduction

Named Entity Recognition is an important pre-processing step for many NLP tasks. As is often the case for NLP tasks, most of the work has been done for English. A major reason for this is the unavailability of manually labeled training data for these languages. In this paper, we address this problem by showing that even in the absence of huge amounts of training data for a given language, unlabeled data from the same language as well as from other languages can be used to improve the existing NER systems. Our approach follows a semi-supervised setting, where, in addition to a small amount of training data, we assume availability of large amounts of unlabeled data from a couple of other languages that are written in a similar or same script as the target language. For example, we hypothesize that data from *German & English* can be used interchangeably to train NER systems for each other.

Our hypothesis stems from the fact that NEs behave similarly across languages (Green et al.,

English	The <b>Obama</b> administration has poured billions of dollars into . . .
German	<b>Barack Obama</b> hat 2012 mit dieser Strategie die Präsidentschaftswahlen gewonnen.
French	. . . et en tentant de faire drailler les nouvelles règles prudentielles, ce démocrate s’est mis dos <b>Barack Obama</b> .

Table 1: Occurrence of **Barack Obama** across three languages.

2011), more so, some of the NEs like the names of the locations and people need not even undergo any orthographic transformation while being used in different languages. For example, *Barack Obama* is spelled the same across all the languages that use roman script like English, German and French (cf. Table 1). We leverage this repetitive information from different languages and show that it can be used to improve the performance of NER system for a given language.

In addition to using manually labeled data for training, we use word clusters obtained from a large monolingual corpus using unsupervised clustering. Word clustering is widely used to reduce the number of parameters in statistical models which leads to improved generalization (Brown et al., 1992; Kneser and Ney, 1993; Koo et al., 2008). Word clusters can effectively capture syntactic, semantic, or distributional regularities among the words belonging to the group (Turian et al., 2010). We acquire such semantic and syntactic similarities from large, unlabeled corpora (§3) that can support the generalization of predictions to new, unseen words in the test set while avoiding overfitting. We show improvements in the NER system performance when informed with these unsupervised clusters for a number of languages (§4.1) as well as from noisy twitter data (§4.2).

## 2 Methodology

Our methodology of using secondary language data to train NER for the target language consists of two steps: (1) Training word clusters from unlabeled secondary language data, (2) Use the word clusters as features in training the target language NER along with labeled target language data.

Since named entities are not a closed word class, it is highly probable that during the test time a named entity is encountered which was not present in the training data. To encounter this sparsity problem, word clusters (trained on a large unlabeled corpus) are used as features in the sequence tagging problems (Clark, 2003; Faruqui and Padó, 2010; Täckström et al., 2012). Thus, an unseen named entity might belong to the same word cluster as some of the seen entities which reinforces the classifier’s belief that it is indeed a named entity, improving the classifier’s performance.

However, the intuition behind using secondary language word clusters as features is that often proper nouns like names of people or locations are spelled the same across orthographically similar languages. Hence, an unseen named entity in the test set might have been present in the word clusters generated from an unlabeled corpus of a secondary language.

## 3 Experimental Setup

**Tools:** We use Stanford Named Entity Recognition system<sup>1</sup> which uses a linear-chain Conditional Random Field to predict the most likely sequence of NE labels (Finkel and Manning, 2009). It uses a variety of features, including the word, lemma, and POS tag of the current word and its context, n-gram features, and word shape. This system supports inclusion of distributional similarity features such as the ones that we want to use in the form of word clusters.

For word clustering, we use the (Clark, 2003) system<sup>2</sup> which in addition to the standard distributional similarity features also uses morphological information about a word using a character-based HMM model for identifying similar words. This gives it the capability to more easily cluster unknown words in morphologically complex lan-

guages like German as compared to only the distributional similarity based approaches (Brown et al., 1992).

**Data:** We evaluate our approach on four different languages namely: *German, English, Spanish & Dutch*. The training and test datasets for German and English were obtained from the shared task “Language Independent Named Entity Recognition”<sup>3</sup> at CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003). The training and test data for Dutch and Spanish were obtained from a similar shared task<sup>4</sup> at CoNLL 2002 (Tjong Kim Sang, 2002). The training data is annotated with four entity types: *person* (PER), *location* (LOC), *organization* (ORG) and *miscellaneous* (MISC).

For generalization data we use the news commentary corpus released by WMT-2012<sup>5</sup> containing articles from 2011. It contains monolingual corpora for *English, German, French, Spanish and Czech*. Each of these corpora on an average have approximately 200 million tokens. Although this corpora collection doesn’t include a corpus for Dutch, we do not search for any other source of Dutch, because our aim in the first place is to show cross-language utility of resources. We train clusters of size 400 for each language as this is suitable number for the size of our generalization corpus (Faruqui and Padó, 2010).

## 4 Results

### 4.1 Training on monolingual clusters

Table 2 shows the  $F_1$  score of the NER systems when trained using the word clusters obtained from five different languages (one at a time). The top row in the table shows the baseline NER system performance trained only on the labeled data without any word clusters. The best results are obtained when word clusters of the same language are used and in all the cases these results are highly statistically significant ( $p < 0.01$ ) according to a McNemar’s test (Dietterich, 1998).

For English, using German word clusters helps improve its performance by 1.3  $F_1$  points which is highly statistically significant. We observe a consistent significant improvement ( $p < 0.05$ ) using

<sup>1</sup><http://nlp.stanford.edu/software/CRF-NER.shtml>

<sup>2</sup><http://www.cs.rhul.ac.uk/home/alexc/pos2.tar.gz>

<sup>3</sup><http://www.cnts.ua.ac.be/conll2003/ner/>

<sup>4</sup><http://www.cnts.ua.ac.be/conll2002/ner/>

<sup>5</sup><http://www.statmt.org/wmt12/translation-task.html>

Word Clusters	English NER	German NER	Dutch NER	Spanish NER
Baseline (None)	91.67	68.27	79.38	79.94
English	<b>93.28**</b>	71.48**	81.65**	80.57*
German	92.92**	<b>77.16**</b>	<b>81.88**</b>	80.4
Spanish	92.43*	70.96**	81.33**	<b>81.57**</b>
French	92.36*	71.04**	81.63**	80.63*
Czech	92.44*	71.47**	80.71**	80.62*
Average	92.69**	72.42**	80.86*	81.34**

Table 2:  $F_1$  scores of NER systems trained across languages. Values in bold indicate highest improvement for that NER. Statistical significance: \*\*  $\rightarrow (p < 0.01)$  and \*  $\rightarrow (p < 0.05)$ .

NER	Baseline	Multilingual	NER	Baseline	Twitter clusters
English	91.67	<b>93.31**</b>	English	91.67	92.61**
German	68.27	<b>77.37**</b>	German	68.27	71.21**
Dutch	79.38	<b>82.15**</b>	Dutch	79.38	80.84**
Spanish	79.94	<b>81.74**</b>	Spanish	79.94	80.39

Table 3:  $F_1$  scores of NER systems trained on multilingual word clusters.

word clusters from French, Spanish and Czech. For a highly morphologically complex language like German it turns out that any amount of data from other languages can be put to good use, as it can be seen that using clusters from any language gives a highly significant improvement in the  $F_1$  score. On an average, an absolute 4 points increase in the score is noted for German.

For Spanish, a significant improvement improvement is obtained when using word clusters from all the languages except for German which might be contributed to the fact they are from different language sub-families. As we do not have generalization data for Dutch, we used word clusters from all other languages and observed that the highest improvement is obtained when word clusters from German are used. This is expected as German and Dutch are typographically very similar languages and come from the Germanic language family. It is notable that even a language like Czech which is a slavic language and is significantly different from the romance and germanic languages gives significant improvement for all the languages.

## 4.2 Training on multilingual clusters

In events when we have unlabeled data from the given language and many secondary languages, we should ideally be able to use word clusters from all of them. In order to do this, we merge the word clusters from different languages together by: (1) Keeping all the words of the given language intact (2) Importing only those words from the secondary language which are not present

Table 4:  $F_1$  scores of NER systems trained using noisy twitter data.

in the original language (thus improving recall). While importing a foreign word, it is assigned to that word cluster which has maximum number of words in common with its present cluster.

Using this technique we merge all the word clusters from different languages (German, English, French, Spanish & Czech) together into one multilingual word clustering. Table 3 shows the performance of multilingual word clusters trained NERs against the baseline NERs. In all the cases, again the NERs trained with multilingual guidance perform significantly better than the baseline NERs and also perform better than the NERs trained with only monolingual word clusters (cf. Table 2).

## 4.3 Training on out-of-domain clusters

The labeled NE training data for all of the languages we have used comes from the newswire domain. Thus the news-commentary data (cf. Sec. 3) that we use for word clustering is in-domain data. Since we cannot always expect to obtain in-domain data, we use word clusters<sup>6</sup> obtained from a large collection of English tweets containing approx. 850 million tokens clustered into 1000 classes for generalization (Owoputi et al., 2013). Table 4 shows the performance of NER systems trained using the twitter word clusters. For English, German and French we again obtain a highly significant improvement; however the improvements obtained using the out-of-domain data are lesser than that obtained using in-domain data.

<sup>6</sup><http://www.ark.cs.cmu.edu/TweetNLP/>

NEs	English	German	Dutch	Spanish	Avg.
LOC	<b>1.9</b>	2.1	<b>4.1</b>	0.9	<b>2.25</b>
MISC	0.9	1.2	0.1	0.4	0.65
ORG	1.7	1.0	3.2	0.1	1.50
PER	1.3	<b>4.6</b>	3.4	<b>1.7</b>	<b>2.75</b>

Table 5: Category-wise improvement ( $\Delta F_1$ ) of NER systems over the baseline model.

English (German)	German (English)	Dutch (German)	Spanish (French)
Schalke	Michoacán	Barrichello	Ledezma
Québec	Bernardin	Israël	Emerson
Lattek	Yalcin	Cofinimmo	Bustamante
Hartmann	Bahraini	Netanyahu	Guayana
Düsseldorf	Nolan	Tamas	Lyon

Table 6: Most frequent out-of-vocabulary words present in the secondary word clusters.

#### 4.4 Analysis

In order to verify our hypothesis that certain NEs, specially the names of *people* and *locations* might not go orthographic changes while transferring to similar languages we look at the category-wise improvement of the NER systems when trained using word clusters of a language (other than self) that gives the best results against the baseline model. For example, for Spanish NER we compare the baseline model with the model trained using French word clusters. Table 5 shows the category-wise improvement in the  $F_1$  score of the NER systems.

For all the languages, the best improvement is obtained for the LOC or the PER class. On an average, the highest improvement is obtained for PER followed by LOC and least for MISC category. The reason for poor improvement in MISC category is that it mostly contains linguistically inflected forms of proper nouns like *Italian*, *Irish*, *Palestinians* etc. which translate into different lexical forms in different languages.

We now analyse the words present in the test set which are out-of-vocabulary (OOV) of the training set. A fraction of these OOV words are present in the word clusters that we obtain from a different language, most of which are names of locations or people as we hypothesised. We list the most frequent such words from the test set in Table 6.

## 5 Related Work

Our work is primarily inspired by Faruqui and Padó (2010) which shows that a substantial improvement in the German NER system performance can be obtained by using unsuper-

vised German word clusters. NER systems have been trained using the same technique for other languages like English (Finkel and Manning, 2009), Croatian (Glavaš et al., 2012) and Slovene (Ljubešić et al., 2012). Other approaches to enhance NER include that of transfer of linguistic structure from one language to another (Täckström et al., 2012; Faruqui and Dyer, 2013) by aligning word clusters across languages. Green et al. (2011) exploits the fact that NEs retain their shape across languages and tries to group NEs across language together.

In a broader perspective this can be framed as a problem of resource sharing (Bateman et al., 2005) among different languages. Languages that are closely related like Hindi and Urdu benefit from sharing resources for NLP tasks (Visweswariah et al., 2010). Also, closely related are the problems of multilingual language learning using unsupervised and supervised approaches (Diab, 2003; Guo and Diab, 2010) and cross-lingual annotation projection applied to bootstrapping parsers (Hwa et al., 2005; Smith and Smith, 2007). Multilingual guidance has also been used for training Part-of-Speech (POS) taggers (Snyder et al., 2008; Snyder et al., 2009; Das and Petrov, 2011).

Our approach is different from the previous approaches in that we are directly using data from secondary languages for training NER systems for the given language instead of deriving any indirect knowledge from the secondary language data using projection or bilingual clustering techniques. It is simple and significantly effective.

## 6 Conclusion

We have shown that a statistically significant improvement in the performance of NER system for a given language can be obtained when the training data is supplemented with word clusters from a secondary language(s) which is written using the same alphabet. The amount of help provided by this secondary language depends on how similar the secondary language is to the given language phylogenetically and also on the domain of the data from which the word clusters are obtained. This performance improvement occurs because many of the NEs, specially, names of *persons* and *locations* remain the same when used in a different language and hence the word class information of such an OOV word is helpful in predicting its NE class.

## References

- [Bateman et al.2005] J. Bateman, I. Kruijff-Korbayov, and G.-J. Kruijff. 2005. Multilingual resource sharing across both related and unrelated languages: An implemented, open-source framework for practical natural language generation. *Research on Language and Computation*, 3(2-3):191–219.
- [Brown et al.1992] P. F. Brown, P. V. deSouza, R. L. Mercer, V. J. D. Pietra, and J. C. Lai. 1992. Class-based n-gram models of natural language. *Comput. Linguist.*, 18(4):467–479, December.
- [Clark2003] A. Clark. 2003. Combining distributional and morphological information for part of speech induction. In *Proceedings of the tenth conference on European chapter of the Association for Computational Linguistics - Volume 1*, EACL '03, pages 59–66, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Das and Petrov2011] D. Das and S. Petrov. 2011. Unsupervised part-of-speech tagging with bilingual graph-based projections. *Proceedings of the Association for Computational Linguistics*.
- [Diab2003] M. T. Diab. 2003. *Word sense disambiguation within a multilingual framework*. Ph.D. thesis, University of Maryland at College Park, College Park, MD, USA. AAI3115805.
- [Dietterich1998] T. G. Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10:1895–1923.
- [Faruqui and Dyer2013] M. Faruqui and C. Dyer. 2013. An information theoretic approach to bilingual word clustering. In *Proceedings of ACL 2013*.
- [Faruqui and Padó2010] M. Faruqui and S. Padó. 2010. Training and Evaluating a German Named Entity Recognizer with Semantic Generalization. In *Proceedings of KONVENS 2010*, Saarbrücken, Germany.
- [Finkel and Manning2009] J. R. Finkel and C. D. Manning. 2009. Nested named entity recognition. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1*, EMNLP '09, pages 141–150, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Glavaš et al.2012] G. Glavaš, M. Karan, F. Šarić, J. Šnajder, J. Mijic, A. Šilic, and B. D. Bašić. 2012. Croner: A state-of-the-art named entity recognition and classification for croatian language. *Money*, 91:94–15.
- [Green et al.2011] S. Green, N. Andrews, M. Gormley, M. Dredze, and C. Manning. 2011. Entity clustering across languages. *NAACL*, pages 591–599.
- [Guo and Diab2010] W. Guo and M. Diab. 2010. Combining orthogonal monolingual and multilingual sources of evidence for all words wsd. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 1542–1551, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Hwa et al.2005] R. Hwa, P. Resnik, A. Weinberg, C. I. Cabezas, and O. Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering*, pages 311–325.
- [Kneser and Ney1993] R. Kneser and H. Ney. 1993. Forming word classes by statistical clustering for statistical language modelling. In R. Khler and B. Rieger, editors, *Contributions to Quantitative Linguistics*, pages 221–226. Springer Netherlands.
- [Koo et al.2008] T. Koo, X. Carreras, and M. Collins. 2008. Simple semi-supervised dependency parsing. In *Proc. of ACL*.
- [Ljubešić et al.2012] N. Ljubešić, M. Stupar, and T. Juric. 2012. Building named entity recognition models for croatian and slovene.
- [Owoputi et al.2013] O. Owoputi, B. OConnor, C. Dyer, K. Gimpel, N. Schneider, and N. A. Smith. 2013. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of NAACL-HLT*, pages 380–390.
- [Smith and Smith2007] D. A. Smith and N. A. Smith. 2007. Probabilistic Models of Nonprojective Dependency Trees. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 132–140, Prague, Czech Republic, June. Association for Computational Linguistics.
- [Snyder et al.2008] B. Snyder, T. Naseem, J. Eisenstein, and R. Barzilay. 2008. Unsupervised multilingual learning for pos tagging. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '08, pages 1041–1050, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Snyder et al.2009] B. Snyder, T. Naseem, J. Eisenstein, and R. Barzilay. 2009. Adding more languages improves unsupervised multilingual part-of-speech tagging: a bayesian non-parametric approach. In *Proceedings NAACL 2009*, pages 83–91, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Täckström et al.2012] O. Täckström, R. McDonald, and J. Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *The 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, page 11. Association for Computational Linguistics.

- [Tjong Kim Sang and De Meulder2003] E. F. Tjong Kim Sang and F. De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In W. Daelemans and M. Osborne, editors, *Proceedings of CoNLL-2003*, pages 142–147. Edmonton, Canada.
- [Tjong Kim Sang2002] E. F. Tjong Kim Sang. 2002. Introduction to the conll-2002 shared task: Language-independent named entity recognition. In *Proceedings of CoNLL-2002*, pages 155–158. Taipei, Taiwan.
- [Turian et al.2010] J. Turian, L. Ratinov, and Y. Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL '10*, pages 384–394, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Visweswariah et al.2010] K. Visweswariah, V. Chenthamarakshan, and N. Kambhatla. 2010. Urdu and hindi: Translation and sharing of linguistic resources. In *Coling 2010: Posters*, pages 1283–1291, Beijing, China, August. Coling 2010 Organizing Committee.